Freeform Search

Database:

US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Term:

Display: |10| **Documents in** Display Format: |-| **Starting with Number** |1|

Generate: ○ **Hit List** ◉ **Hit Count** ○ **Side by Side** ○ **Image**

[Search] [Clear] [Interrupt]

Search History

DATE: Thursday, April 29, 2004   Printable Copy   Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=PGPB,USPT; PLUR=YES; OP=ADJ		
L8	L7 and l6	11	L8
L7	L5 with (((("L3" or level 3 or level three or main or lower) adj3 (cache or memory)))	208	L7
L6	L5 with ((("L2" or level 2 or level two or secondary or second) adj3 cache)	79	L6
L5	buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4)	44160	L5
	DB=USPT; PLUR=YES; OP=ADJ		
L4	(buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4))	32453	L4
L3	l2 and (buffer or flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4)	7	L3
L2	((store-through or writethrough or write-through or write adj through or store adj through) near4 cach$4 and (read adj only or read-only) near4 (instruction adj2 cache) and (store-in or writeback or write-back or write adj back or store adj in) near4 cach$4)	7	L2
	DB=PGPB,USPT; PLUR=YES; OP=ADJ		

h   e b   b   g   e e e   e   fc   f   ff e   ch

L1 (buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4) with (store-in or writeback or write-back or (write adj back) or (store adj in)) with cach$4 and buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4) with ((lower level) or "L3" or (main memory)) and ((read-only or (read adj only)) near4 cach$4))   0   L1

END OF SEARCH HISTORY

Refine Search

Search Results -

Terms	Documents
L13 and ((input or output or first or second) near4 (queue or buffer))	0

Database:

> US Pre-Grant Publication Full-Text Database
> US Patents Full-Text Database
> US OCR Full-Text Database
> EPO Abstracts Database
> JPO Abstracts Database
> Derwent World Patents Index
> IBM Technical Disclosure Bulletins

Search:

> L14

[Refine Search]

[Recall Text] [Clear] [Interrupt]

Search History

DATE: Thursday, April 29, 2004   Printable Copy   Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ			
L14	L13 and ((input or output or first or second) near4 (queue or buffer))	0	L14
L13	L12 and l11	2	L13
L12	L8 with l9	7	L12
L11	L10 with l9	53	L11
L10	((("L3" or level 3 or level three or main or lower) adj3 (cache or memory)))	21596	L10
L9	buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4 or writeback)	8217	L9
L8	(store-in or writeback or write-back or write adj back or store adj in) near4 cach$4	466	L8
DB=PGPB,USPT; PLUR=YES; OP=ADJ			
L7	L6 and ((input or output or first or second) near4 (queue or buffer))	9	L7
L6	L5 and l4	11	L6
L5	L2 with l1	126	L5

h   e b   b cg b   e e ch

L4	L3 with l2	213	L4
L3	((("L3" or level 3 or level three or main or lower) adj3 (cache or memory)))	46319	L3
L2	buffer near6 (flush$4 or cast$3 or castout or purg$4 or empt$4 or remov$4 or replac$4 or eliminat$4 or writeback)	44271	L2
L1	(store-in or writeback or write-back or write adj back or store adj in) near4 cach$4	2398	L1

END OF SEARCH HISTORY

h   e b   b cg b   e e ch

IEEE HOME I SEARCH IEEE I SHOP I WEB ACCOUNT I CONTACT IEEE

◈IEEE

| Membership | Publications/Services | Standards | Conferences | Careers/Jobs |

IEEE Xplore®
RELEASE 1.7

Welcome
United States Patent and Trademark Office

Help   FAQ   Terms   IEEE Peer Review   **Quick Links**

» Se.

Welcome to IEEE Xplore®

- ○ Home
- ○ What Can I Access?
- ○ Log-out

Tables of Contents

- ○ Journals & Magazines
- ○ Conference Proceedings
- ○ Standards

Search

- ○ By Author
- ○ Basic
- ○ Advanced

Member Services

- ○ Join IEEE
- ○ Establish IEEE Web Account
- ○ Access the IEEE Member Digital Library

Your search matched **0** of **1028618** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enteri
new one in the text box.

(flush buffer <sentence> ((store in or writeback or store   **Search**

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine   **CNF** = Conference   **STD** = Standard

Results:
No documents matched your query.

🖶 Print Format

h   eee   e eee   g e ch e ch e   e   e   f h b ffe   e e ce

◈ IEEE

Membership   Publications/Services   Standards   Conferences   Careers/Jobs

IEEE Xplore®
RELEASE 1.7

Welcome
United States Patent and Trademark Office

» Adva

Help   FAQ   Terms   IEEE Peer Review   **Quick Links**

Welcome to IEEE Xplore®

◯ Home
◯ What Can
   I Access?
◯ Log-out

Tables of Contents

◯ Journals
   & Magazines
◯ Conference
   Proceedings
◯ Standards

Search

◯ By Author
◯ Basic
◯ Advanced

Member Services

◯ Join IEEE
◯ Establish IEEE
   Web Account

◯ Access the
   IEEE Member
   Digital Library

Try our New Full-text Search Prototype   **GO**   Help

1) Enter a single keyword, phrase, or Boolean expression.
Example: acoustic imaging (means the phrase acoustic imaging
plus any stem variations)
2) Limit your search by using search operators and field codes,
if desired.
Example: optical <and> (fiber <or> fibre) <in> ti
3) Limit the results by selecting Search Options.
4) Click Search. See Search Examples

```
(flush buffer <sentence>
((store-in or writeback or
store-back) <near/4> cache))
and (flush buffer <sentence>
```

Start Search   **Clear**

Note: This function returns plural and suffixed forms of the
keyword(s).

Search operators: <and> <or> <not> <in> More

Field codes: au (author), ti (title), ab (abstract), jn (publication
name), de (index term) More

Search Options:
Select publication types:
☑ IEEE Journals
☑ IEE Journals
☑ IEEE Conference proceedings
☑ IEE Conference proceedings
☑ IEEE Standards

Select years to search:

From year: |All| to |Present|

Organize search results by:

Sort by: |Relevance|
In: |Descending| order

List |15| Results per page

Home | Log-out | Journals | Conference Proceedings | Standards | Search by Author | Basic Search | Advanced Search | Join IEEE | Web Account |
New this week | OPAC Linking Information | Your Feedback | Technical Support | Email Alerting | No Robots Please | Release Notes | IEEE Online
Publications | Help | FAQ| Terms | Back to Top

h   eee   e eee   g e ch   e ch

● ●

P❦RTAL
(acm)
US Patent & Trademark Office

Search: ◉ The ACM Digital Library   ○ The Guide

(flush buffer <sentence> ((store-in or writeback or store-back| **SEARCH**

▐⁴ F

Terms used
flush **buffer** **sentence** **store** in or **writeback** or **store** **back** **near/4** **cache** and **flush** **buffer** **sentence** **storethrou**

Sort results by [relevant ▾]   ❤ Save results to a Binder
Display results [expanded form ▾]   ? Search Tips
                                                 ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   ne:
Best 200 shown

1   Session S4.1: power in memory and network processors: Embedded cache architecture with
    performance flexibility
    Afzal Malik, Bill Moyer, Roger Zhou
    October 2002   **Proceedings of the international conference on Compilers, architecture, a**
    Full text available: 📄 pdf(122.19 KB)   Additional Information: full citation, abstract, refere

    Next generation portable devices are placing stringent requirements on overall system power and
    wireless internet access are just some of the features being incorporated in these handheld electi
    for high performance and cost sensitive portable products as well as for high end embedded contr
    M¿CORE M2 and M310 fami ...

    Keywords: cache control, cache management, copyback, programmable, push buffer, write buffe

2   An effective write policy for software coherence schemes
    Y.-C. Chen, A. V. Veidenbaum
    December 1992 **Proceedings of the 1992 ACM/IEEE conference on Supercomputing**
    Full text available: 📄 pdf(1.14 MB)   Additional Information: full citation, references, index terms

3   Cache write policies and performance
    Norman P. Jouppi
    May 1993   **ACM SIGARCH Computer Architecture News , Proceedings of the 20th ann**
                    Volume 21 Issue 2
    Full text available: 📄 pdf(1.14 MB)   Additional Information: full citation, abstract, refere

    This paper investigates issues involving writes and caches. First, tradeoffs on writes that miss in tl
    block is fetched on a write miss, whether the missed cache block is allocated in the cache, and wh
    considered. Depending on the combination of these polices chosen, the entire cache miss rate can
    no- ...

4   Coherent network interfaces for fine-grain communication
    Shubhendu S. Mukherjee, Babak Falsafi, Mark D. Hill, David A. Wood
    May 1996   **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd ann**
                    Volume 24 Issue 2
    Full text available: 📄 pdf(1.72 MB)   Additional Information: full citation, abstract, refere

h   c   g e   cf   c

●   ●

Historically, processor accesses to memory-mapped device registers have been marked uncachabl
cache coherence, however, makes it possible for processors and devices to interact with cachable,
performance by facilitating burst transfers of whole cache blocks and reducing control overheads (
interfaces (NIs) that u ...

5 Silent stores for free

Kevin M. Lepak, Mikko H. Lipasti

December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Micr**

Full text available: pdf(521.75 KB) ps(1.97 MB) Additional Information: full citation, references, citings, index
Publisher Site

6 An evaluation of memory consistency models for shared-memory systems with ILP process

Vijay S. Pai, Parthasarathy Ranganathan, Sarita V. Adve, Tracy Harton

September 1996 **Proceedings of the seventh international conference on Architectural sup**
Volume 31 , 30 Issue 9 , 5

Full text available: pdf(1.64 MB) Additional Information: full citation, abstract, referer

Relaxed consistency models have been shown to significantly outperform sequential consistency fc
However, current microprocessors aggressively exploit instruction-level parallelism (ILP) using me
blocking reads. Researchers have conjectured that two techniques, hardware-controlled non-bindii
equalize the hardware pe ...

7 A coherent distributed file cache with directory write-behind

Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart

May 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 2

Full text available: pdf(3.21 MB) Additional Information: full citation, abstract, referer

Extensive caching is a key feature of the Echo distributed file system. Echo client machines maint
write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints on
consistent data structures in the file system even when crashes or network faults prevent some w

Keywords: coherence, file caching, write-behind

8 Cache Memories

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: pdf(4.61 MB) Additional Information: full citation, references, citings, ind

9 Translation lookaside buffer consistency: a software approach

D. L. Black, R. F. Rashid, D. B. Golub, C. R. Hill

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the third inte**
programming languages and operating systems, Volume 17 Issue 2

Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, referer

We discuss the translation lookaside buffer (TLB) consistency problem for multiprocessors, and int
consistency in software. This algorithm has been implemented on several multiprocessors, and is i
basic costs of the algorithm and show that it has minimal impact on application performance. As a
to multiprocessor ...

10
Implementing a cache for a high-performance GaAs microprocessor

h   c   g e   cf   c

O. A. Olukotun, T. N. Mudge, R. B. Brown
April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the 18th ann**
Volume 19 Issue 3
Full text available: pdf(1.12 MB)   Additional Information: full citation, references, citin

11 The directory-based cache coherence protocol for the DASH multiprocessor
Daniel Lenoski, James Laudon, Kourosh Gharachorloo, Anoop Gupta, John Hennessy
May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th ann**
Volume 18 Issue 3
Full text available: pdf(1.74 MB)   Additional Information: full citation, abstract, referer

DASH is a scalable shared-memory multiprocessor currently being developed at Stanford's Compu
processing nodes, each with a portion of the shared-memory, connected to a scalable interconnec
based cache coherence protocol. Unlike traditional snoopy coherence protocols, the DASH protocol
messages sent between th ...

12 Performance of database workloads on shared-memory systems with out-of-order processo
Parthasarathy Ranganathan, Kourosh Gharachorloo, Sarita V. Adve, Luiz André Barroso
October 1998 **Proceedings of the eighth international conference on Architectural suppc**
Volume 33 , 32 Issue 11 , 5
Full text available: pdf(1.62 MB)   Additional Information: full citation, abstract, referer

Database applications such as online transaction processing (OLTP) and decision support systems
market for multiprocessor servers. However, most current system designs have been optimized to
radically different behavior of database workloads (especially OLTP), it is important to re-evaluate
of applicatio ...

13 Memory system performance of programs with intensive heap allocation
Amer Diwan, David Tarditi, Eliot Moss
August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3
Full text available: pdf(2.10 MB)   Additional Information: full citation, abstract, referer

Heap allocation with copying garbage collection is a general storage management technique for pr
performance. To investigate this, we conducted an in-depth study of the memory system performi
machines. We studied the performance of mostly functional Standard ML programs which made he
heap allocation poorly. Howeve ...

Keywords: automatic storage reclamation, copying garbage collection, garbage collection, genera
placement, write through, write-back, write-buffer, write-miss policy, write-policy

14 A scalable approach to thread-level speculation
J. Greggory Steffan, Christopher B. Colohan, Antonia Zhai, Todd C. Mowry
May 2000 **ACM SIGARCH Computer Architecture News , Proceedings of the 27th ann**
Volume 28 Issue 2
Full text available: pdf(186.97 KB)   Additional Information: full citation, abstract, referer

While architects understand how to build cost-effective parallel machines across a wide spectrum (
servers), the real challenge is how to easily create parallel software to effectively exploit all of this
overcoming this problem is Thread-Level Speculation (TLS), which enables the compiler to optimis

15 Compiler and hardware support for cache coherence in large-scale multiprocessors: design
Lynn Choi, Pen-Chung Yew
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd ann**
Volume 24 Issue 2

h   c   g e   cf   c

●   ●

Full text available: 📄 pdf(1.48 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>referei</u>

In this paper, we study a hardware-supported, compiler directed (HSCD) cache coherence scheme off-the-shelf microprocessors, such as the Cray T3D. It can be adapted to various cache organizat architectures. Several system related issues, including critical sections, inter-thread communicatio required hardware sup ...

16 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren
November 2000   **Proceedings of the ninth international conference on Architectural suppor**
28 , 34 Issue 5 , 5
Full text available: 📄 pdf(413.91 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>referei</u>

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-cohe Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiproces four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. ˉ of 32 processors. While s ...

17 Architecture and design of AlphaServer GS320

Kourosh Gharachorloo, Madhu Sharma, Simon Steely, Stephen Van Doren
November 2000   **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: 📄 pdf(1.67 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>referei</u>

This paper describes the architecture and implementation of the AlphaServer GS320, a cache-cohe Compaq. The AlphaServer GS320 architecture is specifically targeted at medium-scale multiproces four Alpha 21264 processors, up to 32GB of coherent memory, and an aggressive IO subsystem. ˉ of 32 processors. While s ...

18 Data speculation support for a chip multiprocessor

Lance Hammond, Mark Willey, Kunle Olukotun
October 1998   **Proceedings of the eighth international conference on Architectural suppo**
Volume 32 , 33 Issue 5 , 11
Full text available: 📄 pdf(1.75 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>referei</u>

Thread-level speculation is a technique that enables parallel execution of sequential applications o implementation of the support for threadlevel speculation on the Hydra chip multiprocessor (CMP) handlers and modifications to the shared secondary cache memory system of the CMP This suppor results show that the s ...

19 Cache behavior of combinator graph reduction

Philip J. Koopman, Peter Lee, Daniel P. Siewiorek
April 1992   **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Vol

Full text available: 📄 pdf(2.18 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>referei</u>

The results of cache-simulation experiments with an abstract machine for reducing combinator gra reduction rates that, for similar kinds of combinator graphs on similar kinds of hardware, compare TIGRE maps easily and efficiently onto standard computer architectures, particularly those that all indication th ...

Keywords: abstract machine, combinators, graph reduction, self-modifying code

20 Improving I/O performance with a conditional store buffer

Lambert Schaelicke, Al Davis
November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Micrc**
Full text available: 📄 pdf(2.53 MB)   Additional Information: <u>full citation</u>, <u>references</u>, <u>index terms</u>

h   c   g   e   cf   c